arXiv:1701.07101v1 [cs.DM] 24 Jan 2017

The switch Markov chain for sampling irregular graphs and digraphs*

Catherine Greenhill†
School of Mathematics and Statistics
UNSW Sydney
NSW 2052, Australia
c.greenhill@unsw.edu.au

Matteo Sfragara‡
Mathematical Institute
Leiden University
P.O. Box 9512, 2300 RA Leiden, The Netherlands
m.sfragara@math.leidenuniv.nl

24 January 2017

Abstract

The problem of efficiently sampling from a set of (undirected, or directed) graphs with a given degree sequence has many applications. One approach to this problem uses a simple Markov chain, which we call the switch chain, to perform the sampling. The switch chain is known to be rapidly mixing for regular degree sequences, both in the undirected and directed setting.

We prove that the switch chain for undirected graphs is rapidly mixing for any degree sequence with minimum degree at least 1 and with maximum degree $d_{\max}$ which satisfies $3 \leq d_{\max} \leq \frac{1}{3}\sqrt{M}$, where M is the sum of the degrees. The mixing time bound obtained is only a factor n larger than that established in the regular case, where n is the number of vertices. Our result covers a wide range of degree sequences, including power-law graphs with parameter $\gamma > 5/2$ and sufficiently many edges.

For directed degree sequences such that the switch chain is irreducible, we prove that the switch chain is rapidly mixing when all in-degrees and out-degrees are positive and bounded above by $\frac{1}{4}\sqrt{m}$, where m is the number of arcs, and not all in-degrees and out-degrees equal 1. The mixing time bound obtained in the directed case is an order of m^2 larger than that established in the regular case.

1 Introduction

There are several approaches to the problem of sampling from a set of graphs (or directed graphs) with a given degree sequence. In this paper we focus on the Markov chain approach. Here the running time of the sampling algorithm must be (deterministically) polynomially

*An earlier version of this work, for undirected graphs only, appeared in SODA 2015 [16].

†Research supported by the Australian Research Council, Discovery Project DP140101519.
‡Research supported by NWO Gravitation Grant 024.002.003–NETWORKS.

bounded but the output need not be exactly uniform: however, the user can specify how far from the uniform distribution the samples may be. Other approaches to the problem of sampling graphs (or directed graphs) are discussed in Section 1.1.

The switch chain is a natural and well-studied Markov chain for sampling from a set of graphs with a given degree sequence. Each move of the switch chain selects two distinct edges uniformly at random and attempts to replace these edges by a perfect matching of the four endvertices, chosen uniformly at random. The proposed move is rejected if the four endvertices are not distinct or if a multiple edge would be formed. We call each such move a *switch*. The precise definitions of the transitions for the switch chain for undirected and directed graphs are given at the start of Sections 2 and 3, respectively.

Ryser [34] used switches to study the structure of 0-1 matrices. Markov chains based on switches have been introduced by Besag and Clifford [5] for 0-1 matrices (bipartite graphs), Diaconis and Sturmfels [9] for contingency tables and Rao, Jana and Bandyopadhyay [33] for directed graphs.

The switch chain is aperiodic and its transition matrix is symmetric. It is well-known that the switch chain is irreducible for any (undirected) degree sequence: see [32, 37]. Irreducibility for the directed chain is not guaranteed, see Rao et al. [33]. However, Berger and Müller-Hanneman [4] and LaMar [24, 26] gave characterisations of directed degree sequences for which the switch chain is irreducible. In particular, the switch chain is irreducible for regular directed graphs (see for example Greenhill [15, Lemma 2.2]).

In order for the switch chain to be useful for sampling, it must converge quickly to its stationary distribution. The rate of convergence of a Markov chain $\mathcal{M}$ is captured by its *mixing time* $\tau(\mathcal{M}, \varepsilon)$, which is the minimum number of steps that the Markov chain $\mathcal{M}$ must run before its distribution is less than ε from stationarity, in total variation distance, from a worst-case starting state. A Markov chain with state space Ω is said to be *rapidly mixing* if its mixing time can be bounded above by some polynomial in $\log(|\Omega|)$ and $\log(\varepsilon^{-1})$. See Section 1.2 for more details.

Cooper, Dyer and Greenhill [7, 8] showed that the switch chain is rapidly mixing for regular undirected graphs. Here the degree $d = d(n)$ may depend on n, the number of vertices. The mixing time bound is given as a polynomial in d and n. Earlier, Kannan, Tetali and Vempala [22] investigated the mixing time of the switch chain for regular bipartite graphs. Greenhill [15] proved that the switch chain for regular directed graphs (that is, d-in, d-out directed graphs) is rapidly mixing, again for any $d = d(n)$. Miklós, Erdős and Soukup [31] proved that the switch chain is rapidly mixing on half-regular bipartite graphs; that is, bipartite degree sequences which are regular for vertices on one side of the bipartition, but need not be regular for the other.

A multicommodity flow argument [35] was used in each of [7, 15, 22, 31] to prove an upper bound on the mixing time of the switch chain. In each case, regularity (or half-regularity) was only required for one lemma, which we will call the *critical lemma*. This is a counting lemma which is used to bound the maximum load of the flow (see [7, Lemma 4], [15, Lemma 5.6] and [31, Lemma 6.15]).

In Section 2 we consider the undirected switch chain and prove the following theorem. This extends the rapid mixing result from [7] to irregular degree sequences which are not too dense.

Given a degree sequence $\boldsymbol{d} = (d_1, \ldots, d_n)$, write $\Omega(\boldsymbol{d})$ for the set of all (simple, undirected) graphs with vertex set $[n] = \{1, 2, \ldots, n\}$ and degree sequence $\boldsymbol{d}$. Recall that $\boldsymbol{d}$ is called *graphical* when $\Omega(\boldsymbol{d})$ is nonempty. We restrict our attention to graphical sequences. Write $d_{\min}$ and $d_{\max}$ for the minimum and maximum degree in $\boldsymbol{d}$, respectively, and let $M = \sum_{j=1}^{n} d_j$ be the sum of the degrees.

Theorem 1.1. *Let $\boldsymbol{d} = (d_1, \ldots, d_n)$ be a graphical degree sequence such that $d_{\min} \geq 1$ and $3 \leq d_{\max} \leq \frac{1}{3}\sqrt{M}$. The mixing time $\tau(\mathcal{M}(\boldsymbol{d}), \varepsilon))$ of the switch chain $\mathcal{M}(\boldsymbol{d})$ with state space $\Omega(\boldsymbol{d})$ satisfies*

$$\tau(\mathcal{M}(\boldsymbol{d}), \varepsilon) \leq d_{\max}^{14} M^9 \left(\tfrac{1}{2} M \log(M) + \log(\varepsilon^{-1})\right).$$

The proof of this result given in an earlier version of this paper [16] had a small gap in the proof. We have fixed the gap here, while also improving the upper bound on $d_{\max}$ by a small constant factor. However, we have not made a serious attempt to optimise the constants.

Theorem 1.1 covers many different degree sequences, for example:

- sparse graphs with constant average degree and maximum degree a sufficiently small constant times $\sqrt{n}$,
- dense graphs with linear average degree and maximum degree a sufficiently small constant times n.
- power-law graphs with parameter $\gamma > 5/2$, when M is sufficiently large. Such graphs were considered by Gao and Wormald [13]: see in particular [13, Section 5], where they prove that $d_{\max} = O(M^{2/5})$ for such graphs (or in their notation, $\Delta = O(M_1^{2/5})$).

Since $M \leq d_{\max} n$, the mixing time bound given above is at most a factor of n larger than that obtained in [7, 8] in the regular case.

The directed case is similar, and is considered in Section 3. To state our main result for the directed switch chain, we must introduce some notation. For definitions about directed graphs not given here, see [1].

A *directed degree sequence* is a sequence $\vec{\boldsymbol{d}}$ of ordered pairs of nonnegative integers $\vec{\boldsymbol{d}} = ((d_1^-, d_1^+), \ldots, (d_n^-, d_n^+))$, such that d_j^- is the in-degree and d_j^+ is the out-degree of vertex j, for all $j \in [n]$. (We use the arrow over the symbol $\vec{\boldsymbol{d}}$ so that our notation distinguishes directed and undirected degree sequences.) The directed degree sequence is *digraphical* if there exists a directed graph with these in-degrees and out-degrees. Write $\Omega(\vec{\boldsymbol{d}})$ for the set of all directed graphs with vertex set $[n]$ such that the in-degree (respectively, out-degree) of vertex j is d_j^- (respectively, d_j^+). Let

$$m = \sum_{j=1}^{n} d_j^- = \sum_{j=1}^{n} d_j^+$$

be the number of arcs in a directed graph with directed degree sequence $\vec{\boldsymbol{d}}$. We say that the directed degree sequence $\vec{\boldsymbol{d}}$ is *switch-irreducible* if the directed switch chain on $\Omega(\vec{\boldsymbol{d}})$ is irreducible.

Finally, let $r_{\min}$ and $r_{\max}$ denote the minimum and maximum *semi-degree* of $\vec{\boldsymbol{d}}$, defined by

$$r_{\min} = \min\{d_1^-,\, d_1^+, \ldots, d_n^-,\, d_n^+\}, \qquad r_{\max} = \max\{d_1^-,\, d_1^+, \ldots, d_n^-,\, d_n^+\}.$$

In [1] these are denoted by $d^0_{\min}$ and $d^0_{\max}$, respectively. However, we prefer the above notation as we must take powers of the maximum semi-degree.

Theorem 1.2. *Let $\vec{\boldsymbol{d}} = ((d_1^-, d_1^+), \ldots, (d_n^-, d_n^+))$ be a digraphical directed degree sequence which is switch-irreducible, such that $r_{\min} \geq 1$ and $2 \leq r_{\max} \leq \frac{1}{4}\sqrt{m}$. Let $\tau(\mathcal{M}(\vec{\boldsymbol{d}}), \varepsilon)$ denote the mixing time of the directed switch chain $\mathcal{M}(\vec{\boldsymbol{d}})$ with state space $\Omega(\vec{\boldsymbol{d}})$. Then*

$$\tau(\mathcal{M}(\vec{\boldsymbol{d}}), \varepsilon) \leq \tfrac{1}{4}\, r_{\max}^{16}\, m^{11}\, \left(m \log(m) + \log(\varepsilon^{-1})\right).$$

Here, since $m \leq r_{\max} n$, the upper bound on $\tau(\mathcal{M}(\vec{\boldsymbol{d}}), \varepsilon)$ given in Theorem 1.2 is an order of m^2 larger than the bound of $50\, d^{25} n^9$ obtained in the directed d-regular case [15], where $r_{\min} = r_{\max} = d$.

The characterisations of Berger and Müller-Hanneman [4] or LaMar [24] can be applied to test whether a given directed degree sequence $\vec{\boldsymbol{d}}$ is switch-irreducible. Note that $r_{\min} \geq 1$ if and only if there are no sources and no sinks (in any directed graph with degree sequence $\vec{\boldsymbol{d}}$). Note that the set of 1-regular directed graphs (with $r_{\min} = r_{\max} = 1$) corresponds to the set of all perfect matchings of $K_{n,n} - M$, where M is a specified perfect matching (forbidding loops in the directed graph). These can be sampled easily, so we assume that $r_{\max} \geq 2$.

We do not believe that the upper bounds given in Theorem 1.1 and Theorem 1.2 are tight. It is likely that the true mixing time in each case is much lower, perhaps $O(dn \log(dn))$ where d is $d_{\max}$ or $r_{\max}$, respectively. Establishing this appears to be far beyond the reach of known proof techniques.

It is not known whether the corresponding counting problems (exact evaluation of $|\Omega(\boldsymbol{d})|$ or $|\Omega(\vec{\boldsymbol{d}})|$) are $\#P$-complete. There are several results giving asymptotic enumeration formulae for $|\Omega(\boldsymbol{d})|$, and some for $|\Omega(\vec{\boldsymbol{d}})|$, under various conditions on the degree sequence: see for example [2, 17, 28, 30] and references therein.

The rest of the paper is structured as follows. In Section 1.1 we review some related results. The necessary Markov chain definitions are given in Section 1.2. Then we consider the undirected switch chain in Section 2, where Theorem 1.1 is proved. Finally, the directed switch chain is studied in Section 3, where Theorem 1.2 is proved.

1.1 History and related work

There are several approaches to the problem of sampling graphs (or directed graphs) with a given degree sequence, though none is known to be efficient for all degree sequences. First we consider undirected graphs. The configuration model of Bollobás [6] gives expected polynomial time uniform sampling if $d_{\max} = O(\sqrt{\log n})$. McKay and Wormald [29] adapted the configuration model to give an algorithm which performs uniform sampling from $\Omega(\boldsymbol{d})$ in expected polynomial time when $d_{\max} = O(M^{1/4})$.

Jerrum and Sinclair [19] used a construction of Tutte's [38] to reduce the problem of approximately sampling from $\Omega(\boldsymbol{d})$ to the problem of approximately sampling perfect matchings from an auxilliary graph. The resulting Markov chain algorithm is rapidly mixing if the degree sequence $\boldsymbol{d}$ is *stable*: see [20]. Stable sequences are those in which small local changes to the degree sequences do not greatly affect the size of $|\Omega(\boldsymbol{d})|$. Specifically, a graphical degree sequence $\boldsymbol{d}$ is stable if

$$(d_{\max} - d_{\min} + 1)^2 \leq 4d_{\min}\,(n - d_{\max} + 1)\,.$$

Many degree sequences which satisfy the conditions of Theorem 1.1 are stable; however, not all stable sequences satisfy the conditions of Theorem 1.1. (For example, if $d_{\min} = n/9$ and $d_{\max} = 4n/9$ then $\boldsymbol{d}$ is stable [20] but $\sqrt{M} \leq 2n/3$, which is not large enough for Theorem 1.1.)

We note that Barvinok and Hartigan [2] showed that the adjacency matrix of a random element of $\Omega(\boldsymbol{d})$ is "close" to a certain "maximum entropy matrix", when the degree sequence is *tame*. The definition of tame depends on the maximum entropy matrix, but a sufficient condition is that $d_{\min} \geq \alpha(n-1)$ and $d_{\max} \leq \beta(n-1)$ for some constants $\alpha, \beta > 0$. Some degree sequences satisfying this latter condition are stable sequences, and many of these degree sequences also satisfy the condition of Theorem 1.1. It would be interesting to explore further the connections between stable degree sequences, tame degree sequences and the mixing rate of the switch Markov chain.

Steger and Wormald [36] gave an easily-implementable algorithm for sampling regular graphs, and proved that their algorithm performs asymptotically uniform sampling in polynomial time when $d = o(n^{1/28})$ (where d denotes the degree). Kim and Vu [23] gave a sharper analysis and established that $d = o(n^{1/3})$ suffices for efficient asymptotically uniform sampling. Bayati, Kim and Saberi [3] extended Steger and Wormald's algorithm to irregular degree sequences, giving polynomial-time asymptotically uniform sampling when $d_{\max} = o(M^{1/4})$. From this they constructed a sequential importance sampling algorithm for $\Omega(\boldsymbol{d})$. A similar approach to that of [29] was described and analysed by Zhao [41] in a general combinatorial setting. Zhao showed that for sampling from $\Omega(\boldsymbol{d})$, when $d_{\max} = o(M^{1/4})$, his algorithm performs asymptotically uniform sampling in time $O(M)$.

There has been less work on the problem of sampling directed graphs with a given degree sequence. However, by characterising directed graphs with as bipartite graphs which avoid a certain perfect matching, it is enough to be able to efficiently sample bipartite graphs with given degrees. The configuration model can be easily adapted for bipartite graphs, and gives expected polynomial time sampling when the product of the maximum in-degree and maximum out-degree is $O(\log n)$, see [27]. Jerrum, Sinclair and Vigoda [21] gave a polynomial-time algorithm for sampling perfect matchings from a given bipartite graph. Combining this with Tutte's construction [38] gives a polynomial-time algorithm for sampling directed graphs with a given degree sequence. Since the auxilliary graph produced by Tutte's construction has a quadratic number of vertices, this method of sampling directed graphs with given degrees has running time bound $O^*(n^{22})$, where $O^*(\cdot)$ ignores logarithmic factors (by [21, Lemma 3.2]). Bayati, Kim and Saberi [3] showed that their sequential importance sampler could be adapted for bipartite graphs with given degrees. Erdős et al. [10] considered an adaptation of the switch chain (which sometimes replaces 3 edges per step, rather than 2) is rapidly mixing

for half-regular bipartite graphs with a forbidden matching, thereby giving an alternative Markov chain for sampling regular directed graphs. (A bipartite graph is half-regular if one vertex bipartition is regular.)

We conclude this section with two recent papers. Erdős, Miklós and Toroczkai [11] showed how to build on the results of [7, 16, 31] using several ingredients including a Markov chain factorisation theorem by the same authors [12] and a certain canonical decomposition of degree sequences due to Tyshkevich [39, 40]. Their approach works by taking degree sequences for which rapid mixing of the switch chain is known, and combining them in order to construct new degree sequences for which they prove that the switch chain is also rapidly mixing. Erdős et al. also considered the directed setting, where Theorem 1.2 now provides a wider range of directed degree sequences for which rapid mixing is known, extending the foundation of the method used in [11]. This should further enlarge the set of more directed degree sequences for which the directed switch chain can be shown to be rapidly mixing.

Gao and Wormald [14] have recently described an extremely efficient expected polynomial time algorithm for exactly uniform sampling d-regular undirected graphs, where $d = o(\sqrt{n})$. The expected running time of their algorithm is $O(d^3 n)$. They also describe a variant of their algorithm with expected running time $O(dn)$ such that the total variation distance of the output distribution from uniform is $o(1)$, again when $d = o(\sqrt{n})$.

1.2 Markov chains and multicommodity flow

For Markov chain definitions not given here, see for example [18].

Let $\mathcal{M}$ be a Markov chain with finite state space Ω, transition matrix P and stationary distribution π. The *total variation distance* between two probability distributions σ, σ' on Ω is given by

$$d_{\mathrm{TV}}(\sigma, \sigma') = \tfrac{1}{2} \sum_{x \in \Omega} |\sigma(x) - \sigma'(x)|.$$

The *mixing time* $\tau(\mathcal{M}, \varepsilon)$ is defined by

$$\tau(\mathcal{M}, \varepsilon) = \max_{x \in \Omega} \min\{T \geq 0 \mid d_{\mathrm{TV}}(P_x^t, \pi) \leq \varepsilon \ \text{ for all } t \geq T\}$$

where P_x^t is the distribution of the state X_t of $\mathcal{M}$ after t steps from the initial state $X_0 = x$.

To bound the mixing time of the switch chain, we apply a multicommodity flow argument. Suppose that $\mathcal{G}$ is the graph underlying a Markov chain $\mathcal{M}$, so that xy is an edge of $\mathcal{G}$ if and only if $P(x, y) > 0$. A *flow* in $\mathcal{G}$ is a function $f : \mathcal{P} \to [0, \infty)$ such that

$$\sum_{p \in \mathcal{P}_{xy}} f(p) = \pi(x)\pi(y) \quad \text{ for all } \ x, y \in \Omega, \ x \neq y.$$

Here $\mathcal{P}_{xy}$ is the set of all simple directed paths from x to y in $\mathcal{G}$ and $\mathcal{P} = \cup_{x \neq y} \mathcal{P}_{xy}$. Extend f to a function on oriented edges by setting $f(e) = \sum_{p \ni e} f(p)$, so that $f(e)$ is the total flow routed through e. Write $Q(e) = \pi(x)P(x, y)$ for the edge $e = xy$. Let $\ell(f)$ be the *length* of the longest path with $f(p) > 0$, and let $\rho(e) = f(e)/Q(e)$ be the *load* of the edge e. The *maximum*

load of the flow is $\rho(f) = \max_e \rho(e)$. Using Sinclair [35, Proposition 1 and Corollary 6'], the mixing time of $\mathcal{M}$ can be bounded above by

$$\tau(\mathcal{M}, \varepsilon) \leq \rho(f)\ell(f) \left(\log(1/\pi^*) + \log(\varepsilon^{-1}\right) \tag{1}$$

where $\pi^* = \min\{\pi(x) \mid x \in \Omega\}$.

2 The undirected switch chain

The (undirected) switch Markov chain $\mathcal{M}(\boldsymbol{d})$ has state space $\Omega(\boldsymbol{d})$ and transitions given by the following procedure: from the current state $G \in \Omega(\boldsymbol{d})$, choose an unordered pair of two distinct non-adjacent edges uniformly at random, say $F = \{\{x, y\}, \{z, w\}\}$, and choose a perfect matching F' from the set of three perfect matchings of (the complete graph on) $\{x, y, z, w\}$, chosen uniformly at random. If $F' \cap (E(G) \setminus F) = \emptyset$ then the next state is the graph G' with edge set $(E(G) \setminus F) \cup F'$, otherwise the next state is $G' = G$.

Define $M_2 = \sum_{j=1}^n d_j(d_j - 1)$. If $P(G, G') \neq 0$ and $G \neq G'$ then $P(G, G') = \frac{1}{3}a(\boldsymbol{d})$, where

$$a(\boldsymbol{d}) = \binom{M/2}{2} - \tfrac{1}{2}\, M_2 \tag{2}$$

is the number of unordered pairs of distinct nonadjacent edges in G. This shows that the switch chain $\mathcal{M}(\boldsymbol{d})$ is symmetric, and it is aperiodic since $P(G, G) \geq 1/3$ for all $G \in \Omega(\boldsymbol{d})$.

In this section we prove Theorem 1.1 by extending the multicommodity flow argument given in [7] in the regular case. The definition of the multicommodity flow given in [7, Section 2.1] carries across to irregular degree sequences without change. This is because the flow from G to G' depends only on the symmetric difference $G \triangle G'$ of G and G', treated as a 2-edge-coloured graph (with edges from $G \setminus G'$ coloured blue and edges from $G' \setminus G$ coloured red, say). The blue degree at a given vertex equals the red degree at that vertex, but in general the blue degree sequence will not be regular. Hence the multicommodity flow definition given in [7] is already general enough to handle irregular degree sequences.

The multicommodity flow is defined using a process which we now sketch. Given $G, G' \in \Omega(\boldsymbol{d})$:

- Define a bijection from the set of blue edges incident at v to the set of red edges incident at v, for each vertex $v \in [n]$. The vector of these bijections is called a *pairing* ψ, and the set of all possible pairings is denoted $\Psi(G, G')$.

- The pairing gives a canonical way to decompose the symmetric difference $G \triangle G'$ into a sequence of *circuits*, where each circuit is a blue/red-alternating closed walk.

- Each circuit is decomposed in a canonical way into a sequence of simpler circuits of two types: 1-*circuits* and 2-*circuits*. A 1-circuit is an alternating cycle in $G \triangle G'$, while a 2-circuit is an alternating walk with one vertex of degree 4, the rest of degree 2, consisting of two odd cycles which share a common vertex. Each 1-circuit or 2-circuit has a designated *start vertex*. (The start vertex of a 2-circuit is the unique vertex of degree 4.) An important fact is that the 1-circuits and 2-circuits are *pairwise edge-disjoint*.

- Each 1-circuit or 2-circuit is processed in a canonical way to give a segment of the canonical path $\gamma_\psi(G, G')$.

For full details see [7, Section 2.1].

2.1 Analysing the flow

Now we show how to bound the load of the flow by adapting the analysis from [7]. Note that some proofs in [7] used the assumption $d = d(n) \leq n/2$, since the general result follows by complementation. This trick does not work for irregular degree sequences, so we cannot make a similar assumption here.

Given matrices G, G', $Z \in \Omega(\boldsymbol{d})$, define the *encoding* L of Z (with respect to G, G') by

$$L + Z = G + G'$$

by identifying each of Z, G and G' with their symmetric 0-1 adjacency matrices. Then L is a symmetric $n \times n$ matrix with entries in $\{-1, 0, 1, 2\}$ and with zero diagonal. Entries which equal -1 or 2 are called *defect entries*. Treating L as an edge-labelled graph with edges labelled $-1, 1, 2$ (and omitting edges corresponding to zero entries), a *defect edge* is an edge labelled -1 or 2. (In [7] these were called "bad edges".) Specifically, we will refer to (-1)-*defect edges* and to 2-*defect edges*. A 2-defect edge is present in both G and G' but is absent in Z, while a (-1)-defect edge is absent in both G and G' but is present in Z.

We say that the *degree* of vertex v in L is the sum of the labels of the edges incident with v (equivalently, the sum of the entries in the row of L corresponding to v). By definition, the degree sequence of L equals $\boldsymbol{d}$.

Some proofs from [7, 8] also apply in the irregular case without any substantial change (after replacing d by $d_{\max}$). These proofs refer only to the symmetric difference and the process used to construct the multicommodity flow (and none of them use the assumption $d \leq n/2$). We state two of these results now.

Lemma 2.1. *Suppose that $G, G', Z, Z' \in \Omega(\boldsymbol{d})$ are such that (Z, Z') is a transition of the switch chain which lies on the canonical path $\gamma_\psi(G, G')$ for some $\psi \in \Psi(G, G')$. Let L be the encoding of Z with respect to (G, G'). Then the following statements hold:*

(i) ([7, Lemma 1]) *From (Z, Z'), L and ψ it is possible to uniquely recover G and G'.*

(ii) ([7, Lemma 2]) *There are at most four defect edges in L. The labelled graph consisting of the defect edges in L must form a subgraph of one of the five possible labelled graphs shown in Figure 1, where "?" represents a label which may be either -1 or 2.*

Proof. The proof of these results for the regular case, [7, Lemma 1, Lemma 2], also applies here since we are using the same multicommodity flow definition, and the proof only uses the symmetric difference of G and G'. (The assumption that $d \leq n/2$, which was sometimes made in [7], is not used in the proof of [7, Lemma 1, Lemma 2].) □

The next result collects together some further useful results about encodings.

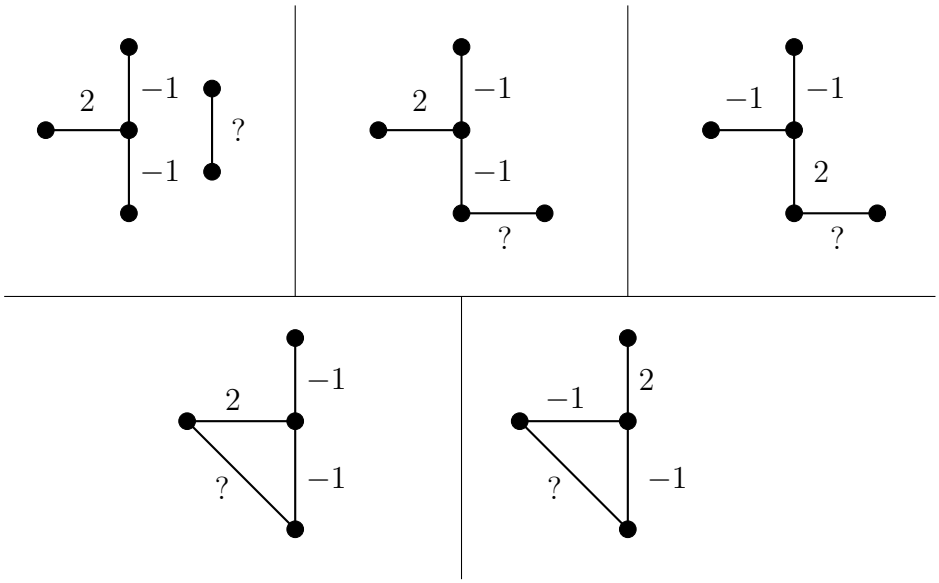


Figure 1: The five possible configurations of four defect edges

Lemma 2.2. *Suppose that the conditions of Lemma 2.1 hold. Let $x, y, z \in [n]$ be distinct vertices.*

(i) *If $L(x, y) = 2$ then $d_x \geq 2$, $d_y \geq 2$.*

(ii) *If $L(x, y) = 2$ and $L(y, z) = 2$ then $d_y \geq 4$.*

(iii) *If $L(x, y) = 2$ and $L(y, z) = -1$ then $d_y \geq 3$.*

Proof. It follows from the definition of the multicommodity flow given in [7] that a 2-defect edge $\{x, y\}$ (with $L(x, y) = 2$) can only arise in two cases:

(a) $\{x, y\}$ is a *shortcut edge* which is present in G, G' but which is absent in Z. (See [7, Figure 4].) In this case, x and y are vertices on some 2-circuit, which is an alternating blue/red walk in the symmetric difference $G \triangle G'$. Hence both x and y have degree at least two in G.

(b) $\{x, y\}$ is an *odd chord* which is present in G, G' but which is absent in Z. (See the section "Processing a 1-circuit" in [7].) In this case, x and y are vertices on some 1-circuit, which is an alternating blue/red walk in the symmetric difference $G \triangle G'$. Again, this shows that both x and y have degree at least two in G.

This proves (i).

Next, if y is incident with two edges of defect 2 then it must be that one is an odd chord for a 1-circuit C_1 and one is a shortcut edge for a 2-circuit C_2, where y does not play the role of x_0 for C_1. Then y is incident in G with an edge of C_1, an edge of C_2 and the two edges

$\{x, y\}$, $\{y, z\}$ which are 2-defect edges in L. Since C_1 and C_2 are edge-disjoint and no defect edge belongs to $G \triangle G'$, it follows that $d_y \geq 4$, proving (ii).

We may adapt this argument to prove (iii), noting that a (-1)-defect edge may only arise from a shortcut edge or an odd chord which is absent in G and G' and present in Z. □

We say that an encoding L is *consistent* with Z if $L + Z$ only takes entries in $\{0, 1, 2\}$. Say that an encoding is *valid* if it satisfies the conclusions of Lemma 2.1(ii), and that a valid encoding is *good* if it also satisfies the conclusion of Lemma 2.2. Let $\mathcal{L}(Z)$ be the set of valid encodings which are consistent with Z, and let $\mathcal{L}^*(Z)$ be the set of good encodings which are consistent with Z. In [7] the set $\mathcal{L}(Z)$ was studied, but we can obtain a slightly better upper bound if we work with the smaller set $\mathcal{L}^*(Z)$.

Lemma 2.3. ([7, Lemma 5] and [8, Lemma 1]) *The load $f(e)$ on the transition $e = (Z, Z')$ satisfies*

$$f(e) \leq d_{\max}^{14} \frac{|\mathcal{L}^*(Z)|}{|\Omega(\boldsymbol{d})|^2}.$$

Proof. In [7, Lemma 5] and [8, Lemma 1] it was shown that $f(e) \leq d_{\max}^{14} |\mathcal{L}(Z)|/|\Omega(\boldsymbol{d})|^2$. (The assumption $d \leq n/2$ is not used in this proof.) The proof relied on the fact that $\mathcal{L}(Z)$ contains all encodings which may arise along a canonical path. But the same is true for $\mathcal{L}^*(Z)$, by Lemma 2.1(ii) and Lemma 2.2, so the proof goes through without change in the irregular setting (after replacing d by $d_{\max}$). □

The switch operation can be extended to encodings in the natural way: each switch reduces two edge labels by one and increases two edge labels by one, without changing the degrees. It was shown in [7, Lemma 3] that from any valid encoding, one could obtain a graph (with no defect edges) by applying a sequence of at most three switches. In [7, Lemma 4] we used this fact to bound the ratio $|\mathcal{L}(Z)|/|\Omega(\boldsymbol{d})|$ for regular degree sequences. This provided an upper bound for the flow $f(e)$ through the transition $e = (Z, Z')$ (as in Lemma 2.3, above). Recall that we now seek an upper bound on the slightly smaller ratio $|\mathcal{L}^*(Z)|/|\Omega(\boldsymbol{d})|$.

The proof of [7, Lemma 3] uses regularity to prove the existence of certain edges exist which are needed in order to find switches to remove the defect edges. This argument fails for irregular degree sequences. However, any argument which gives an upper bound on $|\mathcal{L}^*(Z)|/|\Omega(\boldsymbol{d})|$ will do. So we will instead consider a slightly more complicated operation than a switch, which we call a *3-switch* (this operation is called a "circular C_6-swap" in [10]). (We remark that this approach of considering more complicated operations in order to obtain more freedom has been used to improve asymptotic enumeration results, for example in [30].)

A 3-switch is described by a 6-tuple $(a_1, b_1, a_2, b_2, a_3, b_3)$ of distinct vertices such that $\{a_1, b_1\}$, $\{a_2, b_2\}$, $\{a_3, b_3\}$ are all edges and $\{a_2, b_1\}$, $\{a_3, b_2\}$, $\{a_1, b_3\}$ are all non-edges. The 3-switch deletes the three edges $\{a_1, b_1\}$, $\{a_2, b_2\}$, $\{a_3, b_3\}$ from the edge set and replaces them with $\{a_2, b_1\}$, $\{a_3, b_2\}$, $\{a_1, b_3\}$, as shown in Figure 2.

The 3-switch can also be extended to encodings.

Let $C(p, q)$ be the set of encodings in $\mathcal{L}^*(Z)$ with precisely p defect edges labelled 2 and precisely q defect edges labelled -1, for $p \in \{0, 1, 2\}$ and $q \in \{0, 1, 2, 3\}$. Then $\Omega(\boldsymbol{d}) = C(0, 0)$

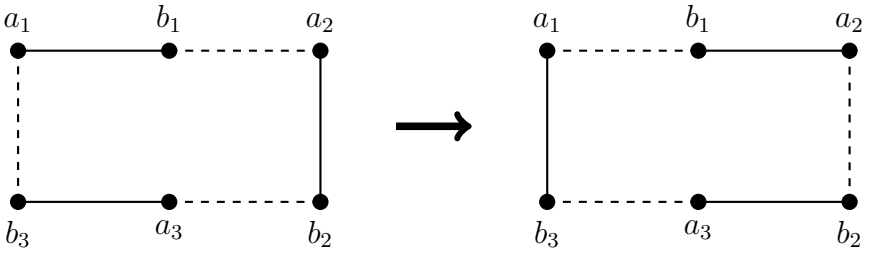


Figure 2: A 3-switch

and

$$\mathcal{L}^*(Z) = \bigcup_{p=0}^{2} \bigcup_{q=0}^{3} C(p,q),$$

where this union is disjoint. (Note that $C(2,3) = \emptyset$, by Lemma 2.1(ii).)

If $L \in C(p,q)$ then there are precisely $M/2 - 2p + q$ non-defect edges in L. (To see this, note that the sum of all entries in the matrix L must equal M, and L has zero diagonal.)

For $v \in [n]$, given an encoding L, write $N_L(v)$ to denote the set of $w \in [n] \setminus \{v\}$ such that $L(v,w) = 1$, and let $\widehat{N}_L(v)$ be the set of all $w \in [n] \setminus \{v\}$ such that $L(v,w) \neq 0$. Then $N_L(v)$ is the set of neighbours of v in L, where neighbours along defect edges are not included, while $\widehat{N}_L(v)$ is the set of all neighbours of v in L (counting both defect and non-defect edges). Let ζ_v (respectively η_v) denote the number of 2-defect edges (respectively, (-1)-defect edges) incident with v in L. Then

$$|N_L(v)| = d_v - 2\zeta_v + \eta_v, \tag{3}$$

and hence the total number of edges incident with v in L is

$$|\widehat{N}_L(v)| = |N_L(v)| + \zeta_v + \eta_v = d_v - \zeta_v + 2\eta_v. \tag{4}$$

With these formulae we can prove the following bounds, which will be very useful in our arguments.

Lemma 2.4. *Suppose that $L \in C(p,q)$ and let a_1, b_1 be distinct vertices with $L(a_1,b_1) \neq 0$.*

(i) *The number of ways to choose an ordered pair of vertices (a_2,b_2) such that $L(a_2,b_2) = 1$ and $L(a_2,b_1) = 0$, with a_1,b_1,a_2,b_2 all distinct, is at least*

$$M - 4p + 2q - \Bigg(d_{\max}\Big(d_{\max} - \zeta_{b_1} + 2\eta_{b_1} + 2\Big) + \eta_{a_1} + \eta_{b_1} - 2(\zeta_{a_1} + \zeta_{b_1}) + \sum_{y \in \widehat{N}_L(b_1)} (\eta_y - 2\zeta_y) \Bigg).$$

(ii) *Now suppose that a_1, b_1, a_2, b_2 are distinct vertices with $L(a_2, b_2) = 1$. Define*

$$\eta^* = \eta_{a_1} + \eta_{b_1} + \eta_{a_2} + \eta_{b_2}, \qquad \zeta^* = \zeta_{a_1} + \zeta_{b_1} + \zeta_{a_2} + \zeta_{b_2}.$$

The number of ways to choose an ordered pair of vertices (a_3, b_3) such that $L(a_3, b_3) = 1$ and $L(a_1, b_3) = L(a_3, b_2) = 0$, with $a_1, b_1, a_2, b_2, a_3, b_3$ all distinct, is at least

$$M - 4p + 2q - \Bigg(d_{\max}\Big(2d_{\max} - (\zeta_{a_1} + \zeta_{b_2}) + 2(\eta_{a_1} + \eta_{b_2}) + 4\Big) + \eta^* - 2\zeta^* + \sum_{x \in \widehat{N}_L(a_1)} (\eta_x - 2\zeta_x) + \sum_{y \in \widehat{N}_L(b_2)} (\eta_y - 2\zeta_y) \Bigg).$$

Proof. For (i), there are $M - 4p + 2q$ possibilities for (a_2, b_2) with $L(a_2, b_2) = 1$, but we must reject the following choices of (a_2, b_2):

- those for which a_1, b_1, a_2, b_2 are not distinct (that is, any choice of (a_2, b_2) which is incident with one of a_1 or b_1), and
- those with $L(a_2, b_1) \neq 0$.

(In [16], we neglected to rule out the possibility that there was a defect edge present between b_1 and a_2. We plug this gap here.) We claim that the number of bad choices for (a_2, b_2) is at most

$$\left(\sum_{y \in \widehat{N}_L(b_1)} |N_L(y)| \right) + |N_L(a_1)| + |N_L(b_1)|. \tag{5}$$

To see this, observe that the sum over y counts all ordered pairs (y, z) with $L(y, z) = 1$ and $L(y, b_1) \neq 0$. This includes each non-defect edge incident with b_1 (when $z = b_1$) and each non-defect edge incident with a_1 (when $y = a_1$). If $L(a_1, b_1) = 1$ then this choice is counted twice, which covers both (a_1, b_1) and (a_2, b_2). Indeed, all edges incident with a_1 or b_1 must be counted twice, to account for the two choices of orientation of $\{a_2, b_2\}$. We achieve this by adding $|N_L(a_1)| + |N_L(b_1)|$ to the upper bound. All other edges counted by the sum over y are of the form (y, z) with

$$y \in \widehat{N}_L(b_1) \setminus \{a_1\}, \qquad z \in N_L(y) \setminus \{b_1\} \quad \text{and} \quad L(y, z) = 1,$$

corresponding a bad choice of $(a_2, b_2) = (y, z)$. This covers all bad choices of (a_2, b_2), completing the proof of (5).

Applying (3) to (5), it follows that the number of bad choices for (a_2, b_2) is at most

$$\left(\sum_{y \in \widehat{N}_L(b_1)} (d_{\max} - 2\zeta_y + \eta_y) \right) + 2d_{\max} - 2(\zeta_{a_1} + \zeta_{b_1}) + \eta_{a_1} + \eta_{b_1}.$$

Applying (4) to b_1 shows that the number of bad choices for (a_2, b_2) is at most

$$d_{\max}\Big(d_{\max} - \zeta_{b_1} + 2\eta_{b_1}\Big) + 2d_{\max} - 2(\zeta_{a_1} + \zeta_{b_1}) + \eta_{a_1} + \eta_{b_1} + \left(\sum_{y \in \widehat{N}_L(b_1)} (\eta_y - 2\zeta_y)\right), \qquad (6)$$

and subtracting this expression from $M - 4p + 2q$ completes the proof of (i).

The bound for (ii) is established in a similar fashion. There are $M - 4p + 2q$ choices for (a_3, b_3) such that $L(a_3, b_3) = 1$. An upper bound on the number of bad choices for (a_3, b_3) is obtained by summing the upper bound for the number of choices which are bad with respect to the pair (a_1, b_1), as given by (6), and the number of choices which are bad with respect to the pair (a_2, b_2). The latter is also given by (6) after replacing a_1 by a_2 and b_1 by b_2, and (for clarity) using the dummy variable x in the sum, rather than y. The proof of (ii) is then completed by subtracting the the sum of these two upper bounds from $M - 4p + 2q$, using the definition of η^*, ζ^*. □

The following lemma is the "critical lemma" which relied on regularity in [7]; its proof is the main task of this section.

Lemma 2.5. *Suppose that $d_{\min} \geq 1$ and $3 \leq d_{\max} \leq \frac{1}{3}\sqrt{M}$. Let $Z \in \Omega(\boldsymbol{d})$. Then*

$$|\mathcal{L}^*(Z)| \leq 2\, M^6\, |\Omega(\boldsymbol{d})|.$$

Proof. We prove that any $L \in \mathcal{L}^*(Z)$ can be transformed into an element of $\Omega(\boldsymbol{d})$ (with no defect edges) using a sequence of at most three 3-switches. The strategy is as follows: in Phase 1 we aim to remove two defects per 3-switch (one 2-defect edge and one (-1)-defect edge), then in Phase 2 we remove one 2-defect edge per 3-switch, and finally in Phase 3 we remove one (-1)-defect edge per 3-switch. There is at most one step in Phase 1, though the other phases may have more than one step: any phase may be empty. Each 3-switch we perform gives rise to an upper bound on certain ratios of the sizes of the sets $C(p, q)$, by double counting. The proof is completed by combining these bounds. (Such an argument is often called a "switching argument" in the asymptotic enumeration literature: see [30] for example.)

Phase 1. If $p + q \leq 3$ then Phase 1 is empty: proceed to Phase 2. Otherwise, suppose that $L \in C(p, q)$ where $p + q = 4$, so $(p, q) \in \{(2, 2), (1, 3)\}$. We count the number of 3-switches $(a_1, b_1, a_2, b_2, a_3, b_3)$ which may be applied to L to produce an encoding $L' \in C_{p-1,q-1}$. This operation is shown in Figure 3, where defect edges are labelled by 2 or -1 and are shown using thicker lines: a thick solid line is a 2-defect edge while a thick dashed line is a (-1)-defect edge.

It follows from Figure 1 that there are at least two choices for a triple of distinct vertices (a_1, b_1, a_2) such that $L(a_1, b_1) = 2$ and $L(a_2, b_1) = -1$.

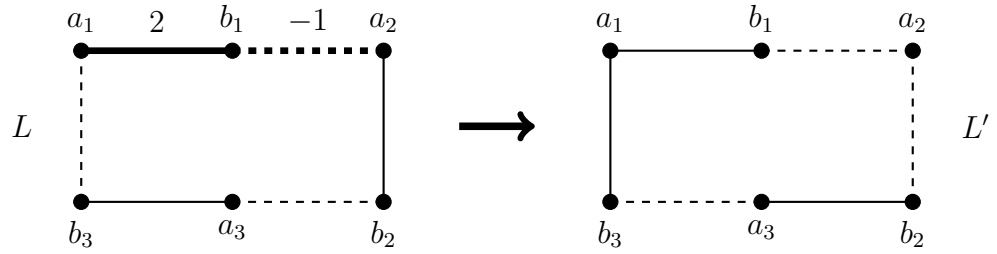


Figure 3: A 3-switch in Phase 1

Given a_1, b_1, a_2, there is at least one vertex $b_2 \in N_L(a_2) \setminus \{a_1\}$. To see this, first suppose that a_2 is not incident with a 2-defect edge. Then $N_L(a_2)$ has at least $d_{a_2} + 1 \geq 2$ elements, leaving at least one which is distinct from a_1. Otherwise, if a_2 is incident with a 2-defect edge then it can be incident with at most one 2-defect edge, since $p \leq 2$. Then there are at least $d_{a_2} - 2$ choices for b_2 in $N_L(a_2) \setminus \{a_1\}$, and this number is positive by Lemma 2.2(iii).

Next, we choose (a_3, b_3) such that all six vertices are distinct, $L(a_3, b_3) = 1$ and $L(a_3, b_2) = L(a_1, b_3) = 0$. We use Lemma 2.4(ii). The worst case is obtained by making $\eta_{a_1} + \eta_{b_2}$ as large as possible, then making η^* as large as possible, while making ζ^* as small as possible. Additionally, adding an edge (with label 1) between some of the known vertices may increase the sum over x or the sum over y. A worst-case example is shown in Figure 4.

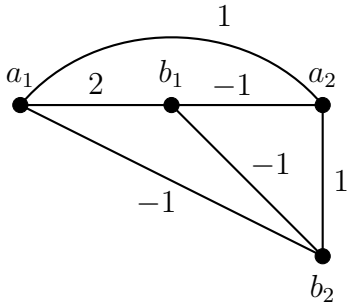


Figure 4: A worst case configuration for the choice of (a_3, b_3) in Phase 1.

Here $(p, q) = (1, 3)$ and

$$\eta^* = 6, \quad \zeta^* = 2, \quad \eta_{a_1} + \eta_{b_2} = 3, \quad \zeta_{a_1} + \zeta_{b_2} = 1,$$
$$\sum_{x \in \widehat{N}_L(a_1)} (\eta_x - 2\zeta_x) = 3, \qquad \sum_{y \in \hat{N}_L(b_2)} (\eta_y - 2\zeta_y) = 0.$$

Substituting these values into Lemma 2.4(ii) shows that the number of good choices for (a_3, b_3) is at least

$$M + 2 - (2d_{\max}^2 + 9d_{\max} + 5) \geq M - 6d_{\max}^2$$

since $d_{\max} \geq 3$.

Combining these estimates shows that the number of possible 3-switches $(a_1, b_1, a_2, b_2, a_3, b_3)$ such that $L(a_1, b_1) = 2$ and $L(a_1, b_3) = -1$ is at least

$$2\left(M - 6d_{\max}^2\right) \geq \frac{2}{3}M, \tag{7}$$

using the fact that $d_{\max} \leq \frac{1}{3}\sqrt{M}$.

Now we consider the reverse of this operation, which is given by reversing the arrow in Figure 3. Given $L' \in C(p-1, q-1)$, we need an upper bound on the number of 6-tuples $(a_1, b_1, a_2, b_2, a_3, b_3)$ such that $L'(a_1, b_1) = L'(a_1, b_3) = L'(a_3, b_2) = 1$ and $L'(a_2, b_1) = L'(a_2, b_2) = L'(a_3, b_3) = 0$. Since the encoding $L \in C(p, q)$ produced by this reverse operation must be consistent with Z, it follows that $\{a_2, b_1\}$ must be an edge of Z. Hence there are at most M choices for (a_2, b_1). Let η'_v denote the number of (-1)-defect edges incident with v in L', for any vertex v. There are at most $d_{b_1} + \eta'_{b_1}$ ways to choose $a_1 \in N_L(b_1)$ and at most $d_{a_1} - 1 + \eta'_{a_1}$ ways to choose $b_3 \in N_L(a_1) \setminus \{b_1\}$. From Figure 1, if $\eta'_{a_1} = 2$ then $\eta'_{b_1} = 0$, and if $\eta'_{b_1} = 1$ then $\eta'_{a_1} \leq 1$. Furthermore, $\eta'_{b_1} \leq 1$. (Otherwise, the reverse switching would produce an encoding which is not valid.) Therefore, the number of ways to choose (a_1, b_3) with the given conditions is at most

$$(d_{b_1} + \eta'_{b_1})(d_{a_1} - 1 + \eta'_{a_1}) \leq d_{\max}\,(d_{\max} + 1) \leq \tfrac{4}{3}\, d_{\max}^2.$$

Finally we must choose (a_3, b_2) such that $L'(a_3, b_2) = 1$, the vertices a_3, b_2 are distinct from the four vertices chosen so far and $L'(a_2, b_2) = L'(a_3, b_3) = 0$. For an upper bound, we simply ensure that (a_3, b_2) is not equal to either orientation of the two edges we have chosen so far (namely (a_1, b_1) or (a_1, b_3) or their reversals). Hence there are at most

$$M - 4(p-1) + 2(q-1) - 4 \leq M$$

good choices for (a_3, b_2). Therefore, the number of ways to apply the reverse operation to $L' \in C(p-1, q-1)$ to produce a consistent encoding $L \in C(p, q)$ is at most $\frac{4}{3}d_{\max}^2 M^2$.

Combining this with (7) shows that whenever $p + q = 4$, by double counting,

$$\frac{|C(p, q)|}{|C(p-1, q-1)|} \leq 2d_{\max}^2 M. \tag{8}$$

Phase 2. Once Phase 1 is complete, we have reached an encoding $L \in C(p, q)$ with $p + q \leq 3$. If $p = 0$ then Phase 2 is empty: proceed to Phase 3. Otherwise, we have

$$(p, q) \in \{(2, 1),\ (2, 0),\ (1, 2),\ (1, 1),\ (1, 0)\}.$$

We count the number of ways to perform a 3-switch to reduce the number of 2-defect edges by one, as shown in Figure 5.

Choose an ordered pair (a_1, b_1) such that $L(a_1, b_1) = 2$, in $2p$ ways. Next, we must choose an ordered pair (a_2, b_2) such that $L(a_2, b_2) = 1$ and $L(a_2, b_1) = 0$ with a_1, b_1, a_2, b_2 all distinct. We will apply Lemma 2.4(i). For a worst case, we make η_{b_1} as large as possible, and then

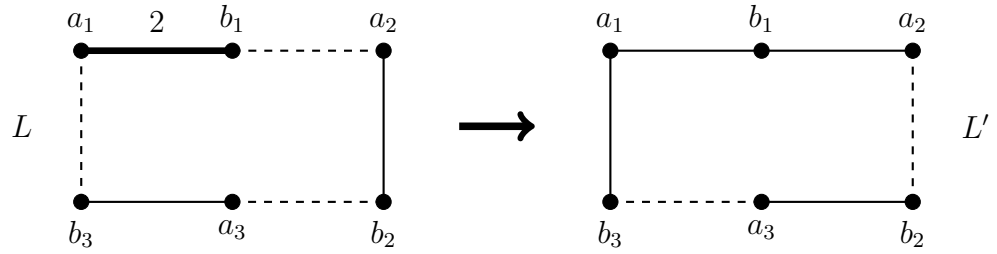


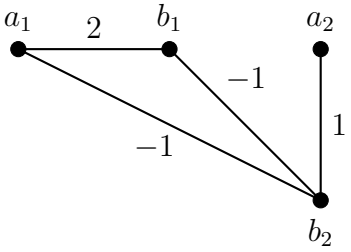
Figure 5: A 3-switch in Phase 2.

make η_{a_1} as large as possible, while making all relevant values of ζ_v as small as possible. The worst case occurs when $(p, q) = (1, 2)$ and $\eta_{b_1} = 2$, which implies that $\eta_{a_1} = 0$. This shows that there are at least

$$M - \left(d_{\max}^2 + 5d_{\max} - 1\right) \geq M - 3d_{\max}^2$$

good choices for (a_2, b_2), using the fact that $d_{\max} \geq 3$.

Next, choose an ordered pair (a_3, b_3) such that all six vertices are distinct, $L(a_3, b_3) = 1$ and $L(a_1, b_3) = L(a_3, b_2) = 0$. We will apply Lemma 2.4(ii). A worst case configuration for (a_3, b_3) is shown below.



Here

$$\eta^* = 4, \quad \zeta^* = 2, \quad \eta_{a_1} + \eta_{b_2} = 3, \quad \zeta_{a_1} + \zeta_{b_2} = 1,$$
$$\sum_{x \in \widehat{N}_L(a_1)} (\eta_x - 2\zeta_x) = 1, \qquad \sum_{y \in \widehat{N}_L(b_2)} (\eta_y - 2\zeta_y) = -2.$$

Plugging these values into the bound from Lemma 2.4(ii), the number of good choices for (a_3, b_3) is at least

$$M - \left(2d_{\max}^2 + 9d_{\max} - 1\right) \geq M - 5d_{\max}^2.$$

Combining these expressions, we conclude that there are at least

$$2\left(M - 3d_{\max}^2\right)\left(M - 5d_{\max}^2\right) \geq \tfrac{16}{27}M^2 \tag{9}$$

valid choices for $(a_1, b_1, a_2, b_2, a_3, b_3)$, using the stated upper bound on $d_{\max}$.

For the reverse operation, let $L' \in C(p-1,q)$ where $p \geq 1$ and $p+q \leq 3$. We need an upper bound on the number of 6-tuples $(a_1,b_1,a_2,b_2,a_3,b_3)$ with $L(a_1,b_1) = L(a_1,b_3) = L(a_2,b_1) = L(a_3,b_2) = 1$ and $L(a_2,b_2) = L(a_3,b_3) = 0$. There are at most

$$M - 4(p-1) + 2q \leq M + 4$$

choices for (a_1,b_1) with $L(a_1,b_1) = 1$, and then there are at most

$$(d_{a_1} - 1 + \eta'_{a_1})(d_{b_1} - 1 + \eta'_{b_1}) \leq d_{\max}^2$$

choices for (a_2,b_3), where (as in Phase 1), η'_v is the number of (-1)-defect edges incident with v in L'. This uses the fact that there are at most two defect edges in L', and hence $\eta'_{a_1} + \eta'_{b_1} \leq 2$, by choice of (a_1,b_1). Finally there are at most

$$M - 4(p-1) + 2q - 6 \leq M - 2$$

valid choices for (a_3,b_2), where for an upper bound we just avoid choosing any orientation of the three edges (namely (a_1,b_1), (a_2,b_1), (a_1,b_3)) which have already been chosen. Hence the number of 6-tuples where the reverse operation can be performed is at most

$$d_{\max}^2(M+4)(M-2) \leq \tfrac{83}{81}\, M^2,$$

since $M \geq 9d_{\max}^2 \geq 81$.

Combining this with (9), it follows that for $(p,q) \in \{(2,1), (2,0), (1,2), (1,1), (1,0)\}$, we have

$$\frac{|C(p,q)|}{|C(p-1,q)|} \leq \tfrac{83}{48}\, d_{\max}^2 < 2d_{\max}^2. \tag{10}$$

Phase 3. After Phase 2, we may suppose that $p = 0$. Let $L \in C(0,q)$ where $q \in \{1,2,3\}$. We count the number of 6-tuples $(a_1,b_1,a_2,b_2,a_3,b_3)$ where a 3-switch can be performed with $L(a_2,b_1) = -1$. Performing this 3-switch will produce $L' \in C(0,q-1)$, as illustrated in Figure 6.

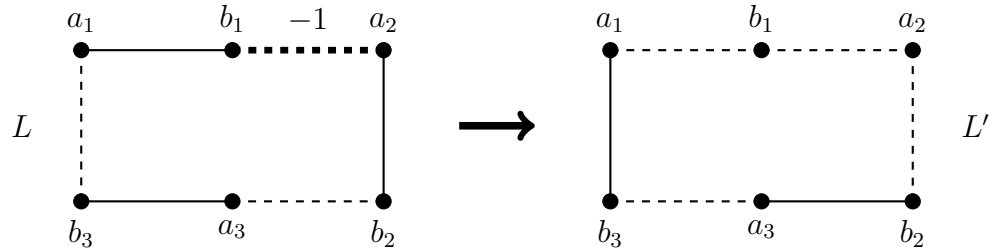


Figure 6: A 3-switch in Phase 3.

There are $2q$ ways to choose (b_1,a_2), and at least $d_{b_1}+1$ ways to choose $a_1 \in N_L(b_1)$. Then there are at least d_{a_2} ways to choose $b_2 \in N_L(a_2) \setminus \{a_1\}$. (Note that the presence of other

(-1)-defect edges incident with b_1 or a_2 only helps here.) Finally, we must choose (a_3, b_3) with $L(a_3, b_3) = 1$ such that all vertices are distinct, $L(a_3, b_2) = 0$ and $L(a_1, b_3) = 0$. Again, an upper bound on the number of bad choices for (a_3, b_3) is given by Lemma 2.4(ii), noting that now $\zeta_v = 0$ for all vertices v. The worst case is attained with $q = 3$, for example when the defect edges are as shown in Figure 4, but with the edge label on $\{a_1, b_1\}$ changed from 2 to 1. Here

$$\eta^* = 6, \quad \eta_{a_1} + \eta_{b_2} = 3, \quad \sum_{x \in \widehat{N}_L(a_1)} \eta_x = 5, \quad \sum_{y \in \widehat{N}_L(b_2)} \eta_y = 4.$$

Substituting these values into Lemma 2.4(ii), the number of valid choices for (a_3, b_3) is at least

$$M - \big(2d_{\max}^2 + 10d_{\max} + 9\big) \geq M - 7d_{\max}^2.$$

Hence the number of 3-switches which can be performed in L to reduce the number of (-1)-defect edges by exactly one is at least

$$2q(d_{b_1} + 1)\, d_{a_2}\, (M - 7d_{\max}^2) \geq 4q(M - 7d_{\max}^2) \geq \tfrac{8}{9}M, \tag{11}$$

using the given bounds on $d_{\min}$ and $d_{\max}$.

For the reverse operation, let $L' \in C(0, q-1)$, where $q \in \{1, 2, 3\}$. We need an upper bound on the number of 6-tuples such that $L(a_1, b_3) = L(a_3, b_2) = 1$, $L(a_1, b_1) = L(b_1, a_2) = L(a_2, b_2) = L(a_3, b_3) = 0$ and $\{a_2, b_1\}$ is an edge of Z. There are at most M choices for (a_2, b_1) satisfying the latter condition, then at most $M + 2(q-1) - 2(d_{a_2} + d_{b_1}) \leq M$ ways to choose (a_3, b_2) with $L(a_3, b_2) = 1$ and a_1, a_3, b_2, b_3 all distinct. (Observe that the presence of any additional (-1)-defect edges incident with a_2 or b_1 can only help here.) Similarly, there at most M ways to choose (a_1, b_3). Hence the number of reverse operations is at most M^3.

Combining this with (11) shows that

$$\frac{|C(0,q)|}{|C(0,q-1)|} \leq \tfrac{9}{8}M^2 \tag{12}$$

for $q \in \{1, 2, 3\}$, by double counting.

Consolidation. Define

$$a = 2d_{\max}^2 M, \qquad b = 2d_{\max}^2, \qquad c = \tfrac{9}{8}M^2.$$

It follows from (8), (10) and (12) that

$$\begin{aligned}\frac{|\mathcal{L}^*(Z)|}{|\Omega(\boldsymbol{d})|} &= \sum_{p=0}^{2}\sum_{q=0}^{3} \frac{|C(p,q)|}{|C(0,0)|} \\ &\leq 1 + b + b^2 + c + bc + b^2c + abc + c^2 + bc^2 + ac^2 + c^3 \\ &\leq 2M^6,\end{aligned}$$

using the upper bound on $d_{\max}$ and the fact that $M \geq 9d_{\max}^2 \geq 81$. This completes the proof of Lemma 2.5. □

Since $M \le d_{\max} n$, the bound $2M^6$ is at most a factor n bigger than the analogous bound $2d^6n^5$ given in [7, Lemma 4] in the regular case.

We can now quickly complete the proof of Theorem 1.1.

Proof of Theorem 1.1. Recall the definitions from Section 1.2. We wish to apply (1). It follows from the configuration model (see [30, Equation (1)]) that the set $\Omega(\boldsymbol{d})$ has size

$$|\Omega(\boldsymbol{d})| \le \frac{M!}{2^{M/2}\,(M/2)!\,\prod_{j=1}^{n} d_j!} \le \exp\left(\tfrac{1}{2}\,M\log(M)\right). \tag{13}$$

Hence the smallest stationary probability π^* satisfies $\log(1/\pi^*) = \log(|\Omega(\boldsymbol{d})|) \le \frac{1}{2}M\log(M)$. Next, $\ell(f) \le M/2$ since each transition along a canonical path replaces an edge of G by an edge of G'.

Finally, if $e = (Z, Z')$ is a transition of the switch chain then $1/Q(e) = 6\,a(\boldsymbol{d}) \le M^2$. Combining this with Lemma 2.3 gives $\rho(f) \le 2d_{\max}^{14}\,M^8$. Substituting these expressions into (1) gives the claimed bound on the mixing time. □

3 The directed switch chain

A directed graph (digraph) $G = (V, A)$ consists of a finite set of vertices and a set A (or $A(G)$) of arcs, where each arc is an ordered pair of distinct vertices. We take $V = [n]$ for some positive integer n.

Recall that $\Omega(\vec{\boldsymbol{d}})$ is the set of all directed graphs with directed degree sequence $\vec{\boldsymbol{d}}$, as defined in Section 1. The directed switch Markov chain, denoted $\mathcal{M}(\vec{\boldsymbol{d}})$, has state space $\Omega(\vec{\boldsymbol{d}})$ and transitions described by the following procedure: from the current digraph $G \in \Omega(\vec{\boldsymbol{d}})$, choose an unordered pair $\{(i,j),(k,\ell)\}$ of distinct arcs of G uniformly at random. If i,j,k,ℓ are distinct and $\{(i,j),(k,\ell)\} \cap A(G) = \emptyset$ then delete the arcs (i,j), (k,ℓ) from G and add the arcs (i,ℓ), (k,j) to obtain the new state; otherwise, remain at G. If distinct digraphs $G, G' \in \Omega(\vec{\boldsymbol{d}})$ are related by a directed switch then $P(G,G') = 1/\binom{m}{2} = P(G',G)$. Hence the directed switch chain $\mathcal{M}(\vec{\boldsymbol{d}})$ is symmetric, so the stationary distribution is uniform over $\Omega(\vec{\boldsymbol{d}})$. The chain is also aperiodic, since for any $G \in \Omega(\vec{\boldsymbol{d}})$ there are at least m pairs of incident edges (for example, pairs of the form $\{(i,j),(j,\ell)\}$). This implies that $P(G,G) \ge m/\binom{m}{2} > 0$. As discussed in Section 1, unlike for undirected graphs, the directed switch chain is not irreducible on $\Omega(\vec{\boldsymbol{d}})$ for all directed degree sequences $\vec{\boldsymbol{d}}$. Instead, we will assume throughout this section that $\vec{\boldsymbol{d}}$ is a switch-irreducible degree sequence; that is, we assume that $\mathcal{M}(\vec{\boldsymbol{d}})$ is irreducible.

In [15], a multicommodity flow analysis was given for $\mathcal{M}(\vec{\boldsymbol{d}})$ for the case of regular directed degree sequences. We now show how to adapt this proof to handle irregular directed degree sequences which satisfy the conditions of Theorem 1.2. One result used to define the multicommodity flow, namely [15, Lemma 2.3], must be reproved here without using the regularity assumption. Having done that, we may use exactly the same multicommodity flow as defined in [15]. This is discussed in Section 3.1 below. Then, the flow is analysed in Section 3.2. As in the undirected case (Section 2), we must reprove a critical counting lemma bounding the number of encodings, without using regularity.

3.1 Defining the flow

The overall structure of the multicommodity flow argument defined in [15] is very similar to the undirected case (on which it was based). Again, given two digraphs $G, G' \in \Omega(\vec{\boldsymbol{d}})$ we consider the symmetric difference $G \triangle G'$ as a 2-arc-coloured digraph. A pairing now consists of a bijection from the blue in-arcs at v to the red in-arcs at v, and a bijection from the blue out-arcs at v to the red out-arcs at v, for each vertex v. With respect to a fixed vertex, the symmetric difference can be decomposed into a sequence of *circuits*, and circuits are then decomposed into 1-circuits or 2-circuits. Both the colour and the direction of the arcs alternate around each 1-circuit or 2-circuit. Then each 1-circuit and 2-circuit must be processed, in order, using a sequence of switches which form part of the canonical path from G to G'.

Some extra cases arise which do not occur in the undirected case, including the special case that the 2-circuit has precisely 6 edges and 3 vertices, which we call a *triangle*: see Figure 7. In order to process a triangle in the regular case, we use some results of LaMar [24, 25] about

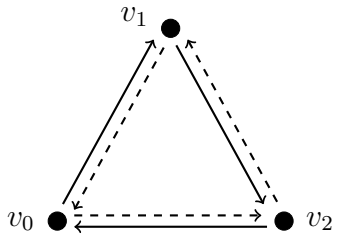


Figure 7: A triangle in the symmetric difference of G and G'

the structure of directed graphs for which the directed switch chain is irreducible. Let (x, U) denote the set of all arcs of the form (x, u) with $u \in U$, and similarly for (U, x). Following LaMar, define the four vertex sets

$$\begin{aligned} U^0 &= \{x \in [n] \setminus U : (x, U) \cup (U, x) \subset A(G)^c\}, \\ U^- &= \{x \in [n] \setminus U : (x, U) \subset A(G) \text{ and } (U, x) \subset A(G)^c\}, \\ U^+ &= \{x \in [n] \setminus U : (x, U) \subset A(G)^c \text{ and } (U, x) \subset A(G)\}, \\ U^\pm &= \{x \in [n] \setminus U : (x, U) \cup (U, x) \subset A(G)\} \end{aligned} \tag{14}$$

where $A(G)^c$ denotes the set of all non-arcs of G.

Combining LaMar's results [24, Theorems 3.3 and 3.4] gives useful characterisation of degree sequences for which $\mathcal{M}(\vec{\boldsymbol{d}})$ is not irreducible which we restate here, for convenience.

Lemma 3.1. [24] *The directed switch chain $\mathcal{M}(\vec{\boldsymbol{d}})$ fails to be irreducible if and only if for every $G \in \Omega(\vec{\boldsymbol{d}})$ there exists a vertex subset $U = \{v_0, v_1, v_2\}$ such that $G[U]$ is a directed 3-cycle and the four sets U^0, U^-, U^+, $U^\pm$ form a partition of $[n] \setminus U$, such that in addition,*

(U1) *no arcs from $U^0 \cup U^+$ to $U^0 \cup U^-$ are present, and*

(U1) *all (non-loop) arcs from $U^- \cup U^\pm$ to $U^+ \cup U^\pm$ are present.*

As in [15], we say that a vertex $x \notin U$ is a *useful neighbour* for the directed 3-cycle on U if $u \notin U^0 \cup U^- \cup U^+ \cup U^{\pm}$. Similarly, we say that the arc (x, y) is a *useful arc* for the directed 3-cycle on U if one of the following conditions holds:

(i) $(x, y) \in A(G)$, with $x \in U^0 \cup U^+$ and $y \in U^0 \cup U^-$, or

(ii) $(x, y) \notin A(G)$, with $x \in U^- \cup U^{\pm}$ and $y \in U^+ \cup U^{\pm}$.

The following lemma is needed in the definition of the multicommodity flow, in order to handle 2-circuits which are triangles. The proof given in the regular case ([15, Lemma 2.3]) used the regularity assumption together with LaMar's characterisation (Lemma 3.1). Below we provide a more direct proof which does not rely on regularity.

Lemma 3.2. *Suppose that $\vec{\boldsymbol{d}}$ is a switch-irreducible directed degree sequence, and that $G \in \Omega(\vec{\boldsymbol{d}})$ contains a set of three vertices $U = \{v_0, v_1, v_2\}$ such that the induced digraph $G[U]$ is a directed 3-cycle. Then there exists a useful neighbour or a useful arc for this 3-cycle.*

Proof. We apply LaMar's characterisation, stated above as Lemma 3.1. Since $\vec{\boldsymbol{d}}$ is switch-irreducible, Lemma 3.1 guarantees that either there exists a vertex $x \in [n] \setminus U$ which does not belong to $U^0 \cup U^- \cup U^+ \cup U^{\pm}$, or there exists an arc (x, y) which contradicts either (U1) or (U2). In the first case x is a useful neighbour of the 3-cycle on U, while in the second case (x, y) is a useful arc for the 3-cycle on U. □

With this lemma in hand, the same definition of multicommodity flow from [15] may be used, for any directed degree sequence $\vec{\boldsymbol{d}}$ which satisfies the conditions of Theorem 1.2.

3.2 Analysing the flow

Given $G, G', Z \in \Omega(\vec{\boldsymbol{d}})$, we define the $n \times n$ matrix L to be an *encoding* of Z (with respect to G, G') by setting $L + Z = G + G'$, as for the undirected case. Treating L as an arc-labelled digraph, we label the arcs with $-1, 1$ or 2 (zero entries correspond to arcs which are absent). A (-1)-*defect arc* is an arc labelled -1. This is an arc which is absent in both G and G' but present in Z. Similarly, a 2-*defect arc* is an arc labelled 2. This is an arc which is present in both G and G' but absent in Z. We write $L(a, b)$ for the label of the directed arc (a, b) in the encoding L.

The next lemma is the directed analogue of Lemma 2.1, collecting together some important results from [15]: the proofs given there did not rely on regularity, and so they extend without change to irregular directed degree sequences.

Lemma 3.3. *Given $G, G' \in \Omega(\vec{\boldsymbol{d}})$ with symmetric difference $G \triangle G'$, let (Z, Z') be a transition on the canonical path from G to G' with respect to the pairing $\psi \in \Psi(G, G')$. Let L be the encoding of Z with respect to (G, G'). Then the following statements hold:*

(i) ([15, Lemma 5.2]) *Given $(Z, Z'), L$ and ψ, there are at most four possibilities for (G, G') such that (Z, Z') is a transition along the canonical path from G to G' corresponding to ψ and L is an encoding for Z with respect to (G, G').*

(ii) ([15, Lemma 5.1]) *There are at most five defect arcs in L. The digraph consisting of the defect arcs in L must form a subdigraph of one of the possible labelled digraphs shown in Figure 8, up to the symmetries described below.*

Define the *arc-reversal* operator ζ, which acts on a digraph G by reversing every arc in G; that is, replacing (u, v) by (v, u) for every arc $(u, v) \in A(G)$. In Figure 8, $\{\mu, \nu\} = \{-1, 2\}$ and $\{\xi, \omega\} = \{-1, 2\}$ independently, giving four symmetries obtained by exchanging these pairs. We can also apply the operation ζ to reverse the orientation of all arcs. Hence each digraph shown in Figure 8 represents up to eight possible digraphs.

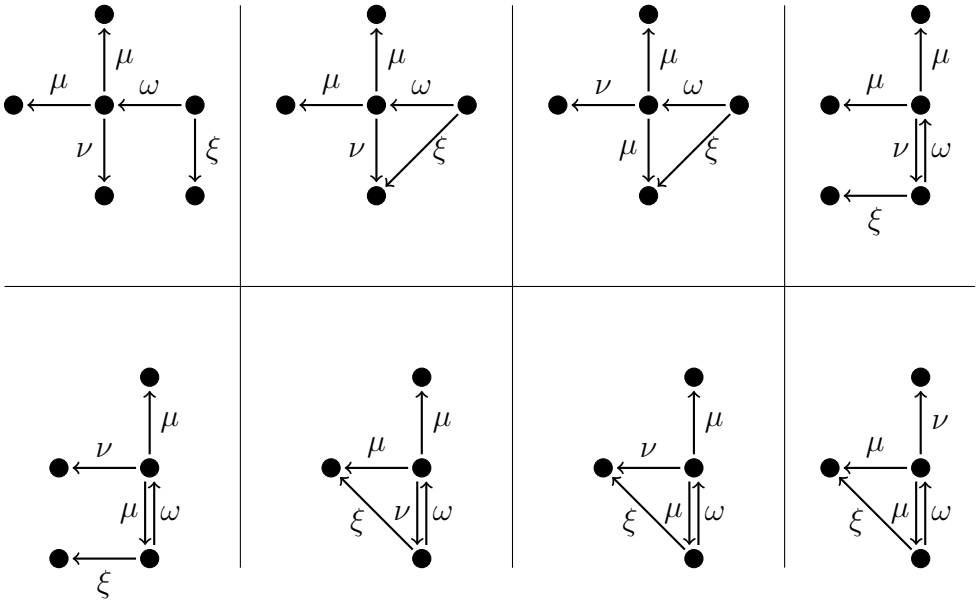


Figure 8: Possible configurations of defect arcs, up to symmetries.

We now generalize the notion of an encoding: given a directed degree sequence $\vec{\boldsymbol{d}}$, an encoding is any $n \times n$ matrix L with entries in $\{-1, 0, 1, 2\}$ such that the jth row sum is the out-degree d_j^+ and the jth column sum is the in-degree d_j^-, for all $j \in [n]$. We say that an encoding L is *consistent* with Z if every entry of $L + Z$ belongs to $\{0, 1, 2\}$, and we say that an encoding L is *valid* if L satisfies the conclusion of Lemma 3.3(ii). Let $\mathcal{L}(Z)$ be the set of valid encodings that are consistent with Z.

The next result is proved just as in the regular case, see [15, Lemma 5.7], since the regularity assumption was not used in the proof given there.

Lemma 3.4. ([15, Lemma 5.7]) *The load $f(e)$ on the transition $e = (Z, Z')$ satisfies*

$$f(e) \leq 4\, r_{\max}^{16}\, \frac{|\mathcal{L}(Z)|}{|\Omega(\vec{\boldsymbol{d}})|^2}.$$

As in Section 2, we can extend the directed switch operation to encodings, ensuring that we never create a label outside the set $\{-1, 0, 1, 2\}$. We wish to use switchings on encodings

to prove an upper bound on the ratio $|\mathcal{L}(Z)|/|\Omega(\vec{\boldsymbol{d}})|$, so we can substitute this bound back into Lemma 3.4. In the regular case, this was achieved using Lemma [15, Lemma 5.5], which proved that from any encoding in $\mathcal{L}(Z)$, one could obtain a digraph in $\Omega(\vec{\boldsymbol{d}})$ using at most three switches. But the proof of this "critical lemma" relied heavily on regularity, so we need a new approach here.

As in the undirected case (Section 2), we introduce a less tightly constrained operation for removing defect edges. A *directed 3-switch* is described by a 6-tuple $(a_1, b_1, a_2, b_2, a_3, b_3)$ of distinct vertices of G such that the arcs (a_1, b_1), (a_2, b_2), (a_3, b_3) are all present in G and the arcs (a_2, b_1), (a_3, b_2), (a_1, b_3) are not. The directed 3-switch deletes the three arcs (a_1, b_1), (a_2, b_2), (a_3, b_3) from the arc set, and replaces them with (a_2, b_1), (a_3, b_2), (a_1, b_3) as shown in Figure 9.

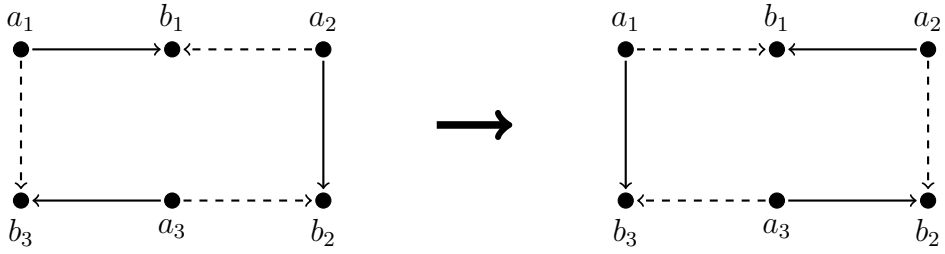


Figure 9: A directed 3-switch.

The directed 3-switch can also be extended to encodings, making sure that after the 3-switch, all labels belong to $\{-1, 0, 1, 2\}$.

Let $C(p, q)$ be the set of encodings in $\mathcal{L}(Z)$ with precisely p defect arcs labelled 2 and precisely q defect arcs labelled -1, for $p, q \in \{0, 1, 2, 3\}$. Then $\Omega(\vec{\boldsymbol{d}}) = C(0, 0)$ and

$$\mathcal{L}(Z) = \bigcup_{p=0}^{3} \bigcup_{q=0}^{3} C(p, q),$$

where this union is disjoint. (Note that $C(3, 3) = \emptyset$ as there are at most five defect arcs, by Lemma 3.3(ii).)

In Section 2 we used a special switch when the maximum number of defect edges were present (that is, in Phase 1): this switch reduced both the number of 2-defect edges and (-1)-defect edges by one. We were able to do this by proving extra structural information about the defect edges (see Lemma 2.2). Unfortunately, we were unable to prove the analogous result in the directed case. (The main difficulty arises from 2-circuits which are triangles.) Hence we will proceed by removing one defect per directed 3-switch, requiring at most 5 directed 3-switches to transform an encoding in $\mathcal{L}(Z)$ into an element of $\Omega(\vec{\boldsymbol{d}})$. Since Phase 1 is missing from our analysis in the directed setting, we rename Phase 2 as Phase A and Phase 3 as Phase B.

If $L \in C(p, q)$ then there are precisely $m - 2p + q$ non-defect arcs in L. Given an encoding L and vertex $v \in [n]$, let $N_L^+(v)$ denote the set of out-neighbours of v (only along non-defect

arcs), that is,

$$N_L^+(v) = \{w \in [n] \setminus \{v\} \mid L(v,w) = 1\}.$$

Similarly,

$$N_L^-(v) = \{w \in [n] \setminus \{v\} \mid L(w,v) = 1\}.$$

We also define the in- and out-neighbourhood when neighbours along defect arcs are included:

$$\widehat{N}_L^+(v) = \{w \in [n] \setminus \{v\} \mid L(v,w) \neq 0\}$$
$$\widehat{N}_L^-(v) = \{w \in [n] \setminus \{v\} \mid L(w,v) \neq 0\}.$$

Recall that the arc (v,w) has *tail* v and *head* w. Let ζ_v^- (respectively, η_v^-) be the number of 2-defect arcs (respectively, (-1)-defect arcs) with v as head, and define ζ_v^+, η_v^+ similarly, for defect arcs with v as tail. The directed analogues of (3) and (4) are

$$|N_L^-(v)| = d_v^- - 2\zeta_v^- + \eta_v^-, \qquad |N_L^+(v)| = d_v^+ - 2\zeta_v^+ + \eta_v^+, \tag{15}$$
$$|\widehat{N}_L^-(v)| = d_v^- - \zeta_v^- + 2\eta_v^-, \qquad |\widehat{N}_L^+(v)| = d_v^+ - \zeta_v^+ + 2\eta_v^+ \ . \tag{16}$$

We can now give the directed analogue of Lemma 2.4.

Lemma 3.5. *Suppose that $L \in C(p,q)$ and let a_1, b_1 be distinct vertices with $L(a_1,b_1) \neq 0$.*

(i) *The number of ways to choose an ordered pair of vertices (a_2,b_2) such that $L(a_2,b_2) = 1$ and $L(a_2,b_1) = 0$, with a_1,b_1,a_2,b_2 all distinct, is at least*

$$m - 2p + q - \Bigg(r_{\max}\Big(r_{\max} - \zeta_{b_1}^- + 2\eta_{b_1}^- + 2\Big) + \eta_{a_1}^- + \eta_{b_1}^+ - 2(\zeta_{a_1}^- + \zeta_{b_1}^+) + \sum_{y \in \widehat{N}_L^-(b_1)} (\eta_y^+ - 2\zeta_y^+) \Bigg).$$

(ii) *Now suppose that a_1, b_1, a_2, b_2 are distinct vertices with $L(a_2,b_2) = 1$. Define*

$$\eta^* = \eta_{a_1}^- + \eta_{b_1}^+ + \eta_{a_2}^- + \eta_{b_2}^+, \qquad \zeta^* = \zeta_{a_1}^- + \zeta_{b_1}^+ + \zeta_{a_2}^- + \zeta_{b_2}^+.$$

The number of ways to choose an ordered pair of vertices (a_3,b_3) such that $L(a_3,b_3) = 1$ and $L(a_1,b_3) = L(a_3,b_2) = 0$, with a_1,b_1,a_2,b_2,a_3,b_3 all distinct, is at least

$$m - 2p + q - \Bigg(r_{\max}\Big(2r_{\max} - (\zeta_{a_1}^+ + \zeta_{b_2}^-) + 2(\eta_{a_1}^+ + \eta_{b_2}^-) + 4\Big) + \eta^* - 2\zeta^* + \sum_{x \in \widehat{N}_L^+(a_1)} (\eta_x^- - 2\zeta_x^-) + \sum_{y \in \widehat{N}_L^-(b_2)} (\eta_y^+ - 2\zeta_y^+) \Bigg).$$

Proof. The proof of (i) follows exactly as in the undirected case. For (ii), an upper bound on the number of bad choices of (a_3, b_3) can be obtained by summing two terms. The first term is the number of choices of (a_3, b_3) such that $\{a_3, b_3\} \cap \{a_2, b_2\} \neq \emptyset$ or $L(a_3, b_2) \neq 0$. An upper bound on this number is given by (i) after replacing a_1 by a_2 and b_1 by b_2. The second term is the number of choices of (a_3, b_3) such that $\{a_3, b_3\} \cap \{a_1, b_1\} \neq \emptyset$ or $L(a_1, b_3) \neq 0$. An upper bound on this number is given by (i) after reversing all arcs: that is, by exchanging the roles of a and b, by exchanging the superscripts "+" and "−", and (for clarity) replacing the dummy variable y by x. The proof of (ii) is completed by adding these two terms together and subtracting them from $m - 2p + q$, using the definition of η^* and ζ^*. □

The following lemma is the "critical lemma" for the directed case, which we prove in the irregular setting by adapting the proof of Lemma 2.5.

Lemma 3.6. *Suppose that the directed degree sequence $\vec{\boldsymbol{d}}$ satisfies $r_{\min} \geq 1$ and $2 \leq r_{\max} \leq \frac{1}{4}\sqrt{m}$. Let $Z \in \Omega(\vec{\boldsymbol{d}})$. Then*

$$|\mathcal{L}(Z)| \leq \tfrac{1}{8} m^8 \, |\Omega(\vec{\boldsymbol{d}})|.$$

Proof. We prove that any $L \in \mathcal{L}(Z)$ can be transformed into an element of $\Omega(\vec{\boldsymbol{d}})$ (with no defect arcs) using a sequence of at most five directed 3-switches. An A-switch reduces the number of 2-defect arcs by one, without changing the number of (-1)-defect arcs. A B-switch takes an encoding with no 2-defect arcs, and reduces the number of (-1)-defect arcs by one. Given an encoding $L \in C(p, q)$, in Phase A we perform an A-switch p times, giving an encoding with no 2-defect arcs. Then in Phase B we perform a B-switch q times, reaching an element of $\Omega(\vec{\boldsymbol{d}})$ (that is, an encoding with no defect arcs).

Phase A. If $p = 0$ then Phase A is empty and we proceed to Phase B. Now assume that $p \in \{1, 2, 3\}$ and let $L \in C(p, q)$, where $q \in \{0, 1, 2, 3\}$ and $p + q \leq 5$. We want a lower bound on the number of 6-tuples $(a_1, b_1, a_2, b_2, a_3, b_3)$ where a 3-switch can be performed with $L(a_1, b_1) = 2$: performing this switch we will produce $L' \in C(p-1, q)$, as shown in Figure 10. There are p ways to choose the 2-defect arc (a_1, b_1). Next, the number of choices of (a_2, b_2)

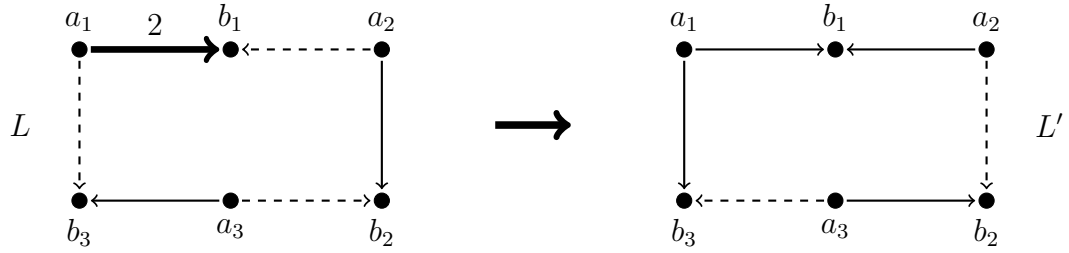


Figure 10: An A-switch.

such that a_1, b_1, a_2, b_2 are all distinct, $L(a_2, b_2) = 1$ and $L(a_2, b_1) = 0$ given by Lemma 3.5(ii). A worst case configuration for the choice of (a_2, b_2) is shown in Figure 11. Here $\eta^-_{b_1} = 2$, $\eta^-_{a_1} + \eta^+_{b_1} = 2$ and $\sum_{y \in \widehat{N}^-_L(b_1)} \eta^+_y = 2$. (Note that we need only consider configurations of

defect arcs which are subdigraphs of those shown in Figure 8: in particular, $\eta_v^- \leq 2$ for all v.) Clearly $q = 3$, and it is possible that $p = 2$, though for an upper bound the additional 2-defect arc will not be incident with any of a_1, b_1, a_2, b_2. Substituting these values into Lemma 3.5(i) shows that the number of valid choices for (a_2, b_2) is at least

$$m - 1 - (r_{\max}^2 + 5r_{\max} + 2) \geq m - 5r_{\max}^2,$$

using the fact that $r_{\max} \geq 2$.

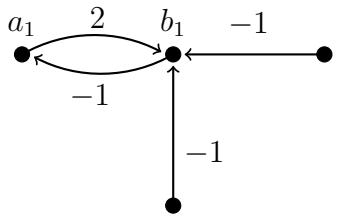
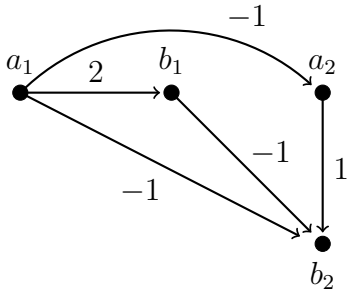


Figure 11: A worst-case configuration for the choice of (a_2, b_2) in Phase A.

Finally we must choose (a_3, b_3) such that the six chosen vertices are distinct, $L(a_3, b_3) = 1$ and $L(a_3, b_2) = L(a_1, b_3) = 0$. We obtain an upper bound on the number of valid choices for (a_3, b_3) by applying Lemma 3.5(ii). A worst case configuration for the choice of (a_3, b_3) is shown below.



Here

$$\eta^* = 2, \quad \zeta^* = 0, \quad \eta_{a_1}^+ + \eta_{b_2}^- = 4, \quad \zeta_{a_1}^+ + \zeta_{b_2}^- = 1,$$
$$\sum_{x \in \widehat{N}_L^+(a_1)} (\eta_x^- - 2\zeta_x^-) = 1, \qquad \sum_{y \in \widehat{N}_L^-(b_2)} (\eta_y^+ - 2\zeta_y^+) = 2.$$

Clearly $q = 3$, and it is possible that $p = 2$, though for a worst case the remaining 2-defect arc will not be incident with any of the vertices shown. Substiting these values into Lemma 3.5(ii), we find that the number of valid choices for (a_3, b_3) is at least

$$m - 1 - (2r_{\max}^2 + 11r_{\max} + 5) \leq m - 9r_{\max}^2.$$

Therefore, using the upper bound on $r_{\max}$, there are at least

$$p\,(m - 5r_{\max}^2)(m - 9r_{\max}^2) \geq \tfrac{77}{256}m^2 > \tfrac{3}{10}m^2 \tag{17}$$

choices of A-switch which can be performed in L to give an element of $C(p-1,q)$.

Now we consider the reverse operation. Let $L' \in C(p-1,q)$, where $p \in \{1,2,3\}$ and $q \in \{0,1,2,3\}$. We want an upper bound on the number of 6-tuples $(a_1,b_1,a_2,b_2,a_3,b_3)$ such that $L'(a_1,b_1) = L'(a_2,b_1) = L'(a_3,b_2) = L'(a_1,b_3) = 1$ and $L'(a_2,b_2) = L'(a_3,b_3) = 0$. There are at most $m-2(p-1)+q \le m+3$ choices for the arc (a_1,b_1) with $L'(a_1,b_1) = 1$, and then there are at most

$$(d^+_{a_1} + (\eta^+_{a_1})' - 1)(d^-_{b_1} + (\eta^-_{b_1})' - 1) \le r_{\max}(r_{\max}+1) \le \tfrac{3}{2} r^2_{\max}$$

choices for the ordered pair (a_2,b_3), since $r_{\max} \ge 2$. Here $(\eta^-_{b_1})'$, respectively $(\eta^+_{a_1})'$, denotes the number of (-1)-defect arcs in L' with b_1 as head (respectively, with a_1 as tail). The worst case is when $q=3$ and $\{(\eta^+_{a_1})',\, (\eta^-_{b_1})'\} = \{1,2\}$. (Note that $L'(a_1,b_1) = 1$, so there can be no (-1)-defect arc from a_1 to b_1.)

Finally, there are at most $m-2(p-1)+q-3 \le m$ choices for the arc (a_3,b_2), since this arc must be distinct from the 3 arcs chosen so far. Hence, the number of 6-tuples where the reverse operation can be performed in L' is at most

$$\tfrac{3}{2} r^2_{\max}\, m(m+3) \le \tfrac{201}{128} r^2_{\max}\, m^2 < \tfrac{8}{5}\, r^2_{\max}\, m^2,$$

using the fact that $m \ge 16 r^2_{\max} \ge 64$.

Combining this upper bound with (17) gives, by double counting,

$$\frac{|C(p,q)|}{|C(p-1,q)|} \le \tfrac{16}{3} r^2_{\max}. \tag{18}$$

Phase B. We proceed with Phase B when there are no 2-defect arcs in our encoding.

Let $L \in C(0,q)$ with $q \in \{1,2,3\}$. We want a lower bound to the number of 6-tuples $(a_1,b_1,a_2,b_2,a_3,b_3)$ where a 3-switch can be performed in L to produce an encoding $L' \in C(0,q-1)$, as shown in Figure 12.

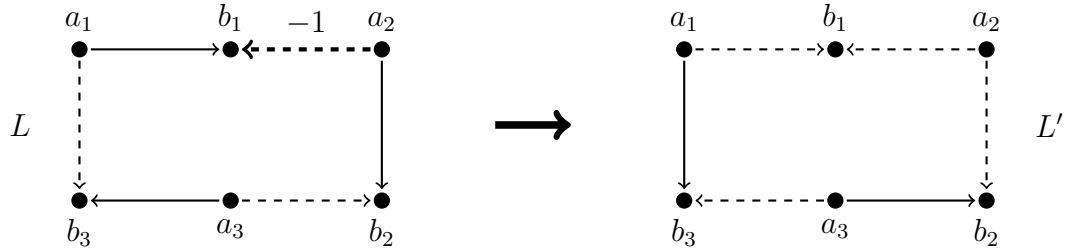


Figure 12: A B-switch.

There are exactly $q \ge 1$ choices for the (-1)-defect arc (a_2,b_1). Then we can choose $a_1 \in N^-_L(b_1)$ in $d^-_{b_1} + \eta^-_{b_1} \ge 2$ ways, and we can choose $b_2 \in N^+_L(a_2) \setminus \{a_1\}$ in $d^+_{a_2} + \eta^+_{a_2} - 1 \ge 1$ way. It remains to choose (a_3,b_3) with all six chosen vertices distinct, $L(a_3,b_3) = 1$ and $L(a_3,b_2) = L(a_1,b_3) = 0$. A lower bound for the number of valid choices of (a_3,b_3) is given by Lemma 3.5(ii). A worst case configuration for the choice of (a_3,b_3) is shown in Figure 13.

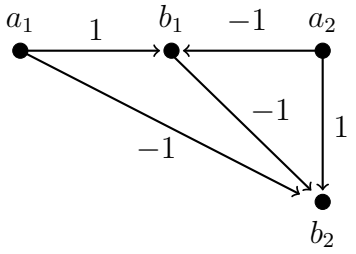


Figure 13: A worst case configuration for the choice of (a_3, b_3) in Phase B.

Here $\eta^+_{a_1} + \eta^-_{b_2} = 3$ and $\eta^* = 1$, with $\sum_{x \in \widehat{N}^+_L(a_1)} \eta^-_x = \sum_{y \in \widehat{N}^-_L(b_2)} \eta^+_y = 3$. Hence the number of valid choices for (a_3, b_3) is at least

$$m + 3 - (2r^2_{\max} + 10r_{\max} + 7) \geq m - 8r^2_{\max},$$

since $r_{\max} \geq 2$.

Putting this together, the number of 6-tuples $(a_1, b_1, a_2, b_2, a_3, b_3)$ where a B-switch can be performed in L to give an element of $C(0, q-1)$ is at least

$$2q\,(m - 8r^2_{\max}) \geq m, \tag{19}$$

using the stated upper bound on $r_{\max}$.

Next we consider the reverse operation. Let $L' \in C(0, q-1)$, where $q \in \{1, 2, 3\}$. We want an upper bound on the number of 6-tuples with $L'(a_3, b_2) = L'(a_1, b_3) = 1$ and $L'(a_1, b_1) = L'(a_2, b_1) = L'(a_2, b_2) = L'(a_3, b_3) = 0$. The encoding L produced by the reverse operation must be consistent with Z after the reverse operation. Since $L(a_2, b_1) = -1$, this implies that (a_2, b_1) must be an arc of Z. Hence there are at most m choices for (a_2, b_1). Next we choose (a_1, b_3) with $L(a_1, b_3) = 1$. For an upper bound, we only check that the arc (a_1, b_3) is not incident with b_1 or a_2, ruling out at least

$$d^-_{a_2} + d^+_{a_2} + d^-_{b_1} + d^+_{b_1} - 2 \geq 2$$

choices. (The -2 term avoids any double-counting.) Therefore there are at most

$$m + (q-1) - 2 \leq m$$

choices for (a_1, b_3). The same argument shows that there are at most m choices for the arc (a_3, b_2). It follows that the number of 6-tuples where the reverse operation can be performed in L' to produce an element of $C(0, q)$ is at most m^3.

Combining this bound with (19) gives, by double counting,

$$\frac{|C(0, q)|}{|C(0, q-1)|} \leq m^2. \tag{20}$$

Consolidation. Define

$$a = \tfrac{16}{3} r^2_{\max}, \qquad b = m^2.$$

Then, recalling that $|C(3,3)| = 0$,

$$\frac{|\mathcal{L}(Z)|}{|\Omega(\vec{\boldsymbol{d}})|} = \sum_{p=0}^{3}\sum_{q=0}^{3}\frac{|C(p,q)|}{|C(0,0)|} \leq (1+a+a^2+a^3)(1+b+b^2+b^3) - a^3b^3$$
$$\leq \tfrac{1}{8}m^8.$$

The last inequality follows since $r_{\max} \geq 2$ amd $m \geq 16r_{\max}^2 \geq 64$. This completes the proof of Lemma 3.6. □

Now we can prove our upper bound on the mixing time of the directed switch chain.

Proof of Theorem 1.2. Recall the definitions from Section 1.2. It follows from the bipartite model of directed graphs that

$$|\Omega(\vec{\boldsymbol{d}})| \leq m! \leq \sqrt{2\pi m}\,\left(\frac{m}{e}\right)^m.$$

Therefore the smallest stationary probability π^* satisfies $\log(1/\pi*) = \log|\Omega(\vec{\boldsymbol{d}})| \leq m\log m$. Next, observe that $\ell(f) \leq m$ since each transition along a canonical path replaces an edge of G by an edge of G'. Finally, if $e = (Z, Z')$ is a transition of the directed switch chain then $1/Q(e) = \binom{m}{2}|\Omega(\vec{\boldsymbol{d}})|$. Combining this with Lemmas 3.4 and 3.6 gives $\rho(f) \leq \frac{1}{4}r_{\max}^{16}m^{10}$. Substituting these expressions into (1) completes the proof. □

References

[1] J. Bang-Jensen and G. Gutin, *Digraphs: Theory, Algorithms and Applications (2nd edn.)*, Springer, New York, 2008.

[2] A. Barvinok and J.A. Hartigan, The number of graphs and a random graph with a given degree sequence, *Random Structures and Algorithms* **42** (2013), 301–348.

[3] M. Bayati, J.H. Kim and A. Saberi, A sequential algorithm for generating random graphs, *Algorithmica* **58** (2010), 860–910.

[4] A. Berger and M. Müller–Hannemann, Uniform sampling of digraphs with a fixed degree sequence, in *Graph Theoretic Concepts in Computer Science*, Lecture Notes in Computer Science vol. 6410, Springer, Berlin, 2010, pp. 220–231.

[5] J. Besag and P. Clifford, Generalized Monte Carlo significance tests, *Biometrika* **76** (1989), 633–642.

[6] B. Bollobás, A probabilistic proof of an asymptotic formula for the number of labelled regular graphs, *European Journal of Combinatorics* **1** (1980), 311–316.

[7] C. Cooper, M.E. Dyer and C. Greenhill, Sampling regular graphs and a peer-to-peer network, *Combinatorics, Probability and Computing* **16** (2007), 557–593.

[8] C. Cooper, M.E. Dyer and C. Greenhill, Corrigendum: Sampling regular graphs and a peer-to-peer network. arXiv:1203.6111

[9] P. Diaconis and B. Sturmfels, Algebraic algorithms for sampling from conditional distributions, The Annals of Statistics **26** (1998), 363–397.

[10] P.L. Erdős, S.Z. Kiss, I. Miklós and L. Soukup, Approximate counting of graphical realizations, *PLoS ONE* **10** (2015), e0131300.

[11] P.L. Erdős, I. Miklós and Z. Toroczkai, New classes of degree sequences with fast mixing swap Markov chain sampling, Preprint (2016). arXiv:1601.08224

[12] P.L. Erdős, I. Miklós and Z. Toroczkai, A decomposition based proof for fast mixing of a Markov chain over balanced realizations of a joint degree matrix, *SIAM Journal of Discrete Mathematics* **29** (2015), 481–499.

[13] P. Gao and N. Wormald, Enumeration of graphs with a heavy-tailed degree sequence, Preprint, 2014. arXiv:1404.1250

[14] P. Gao and N. Wormald, Uniform generation of random regular graphs, Preprint, 2015. arXiv:1511.01175

[15] C. Greenhill, A polynomial bound on the mixing time of a Markov chain for sampling regular directed graphs, *Electronic Journal of Combinatorics* **18** (2011), #P234.

[16] C. Greenhill, The switch Markov chain for sampling irregular graphs (extended abstract), in *Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms*, New York–Philadelphia, 2015, pp. 1564–1572.

[17] C. Greenhill and B. D. McKay, Random dense bipartite graphs and directed graphs with specified degrees, *Random Structures and Algorithms* **35** (2009), 222–249.

[18] M. Jerrum, *Counting, sampling and integrating: algorithms and complexity*, Lectures in Mathematics – ETH Zürich, Birkhäuser, Basel, 2003.

[19] M. Jerrum and A. Sinclair, Fast uniform generation of regular graphs, *Theoretical Computer Science* **73** (1990), 91–100.

[20] M. Jerrum, A. Sinclair and B. McKay, When is a graphical sequence stable?, in *Random Graphs, Vol. 2 (Poznán, 1989)*, Wiley, New York, 1992, pp. 101–115.

[21] M. Jerrum, A. Sinclair and E. Vigoda, A polynomial-time approximation algorithm for the permanent of a matrix with non-negative entries, *Journal of the Association for Computing Machinery* **51** (2004), 671–697.

[22] R. Kannan, P. Tetali and S. Vempala, Simple Markov-chain algorithms for generating bipartite graphs and tournaments, *Random Structures and Algorithms* **14** (1999), 293–308.

[23] J.H. Kim and V. Vu, Generating random regular graphs, *Combinatorica* **26** (2006), 683–708.

[24] M.D. LaMar, On uniform sampling simple directed graph realizations of degree sequences, Preprint (2009). arXiv:0912:3834

[25] M.D. LaMar, Algorithms for realizing degree sequences of directed graphs, Preprint (2009). arXiv:0906.0343

[26] M.D. LaMar, Directed 3-cycle anchored digraphs and their application in the uniform sampling of realizations from a fixed degree sequence, in *Proceedings of the 2011 Winter Simulation Conference (WSC)*, IEEE Press, New York, 2011.

[27] B.D. McKay, Asymptotics for 0-1 matrices with prescribed line sums, *Enumeration and Design*, Academic Press, Toronto, 1984, pp. 225–238.

[28] B.D. McKay, Subgraphs of dense random graphs with specified degrees, *Combinatorics, Probability and Computing* **20** (2011), 413–433.

[29] B.D. McKay and N.C. Wormald, Uniform generation of random regular graphs of moderate degree, *Journal of Algorithms* **11** (1990), 52–67.

[30] B.D. McKay and N.C. Wormald, Asymptotic enumeration by degree sequence of graphs with degrees $o(n^{1/2})$, *Combinatorica* **11** (1991), 369–382.

[31] I. Miklós, P.L. Erdős and L. Soukup, Towards random uniform sampling of bipartite graphs with given degree sequence, *Electronic Journal of Combinatorics* **20(1)**, 2013, #P16.

[32] J. Petersen, Die Theorie der regulären graphs, *Acta Mathematica* **15** (1891), 193–220.

[33] A.R. Rao, R. Jana and S. Bandyopadhyay, A Markov chain Monte Carlo method for generating random (0,1)-matrices with given marginals, *Sankhyā: The Indian Journal of Statistics* **58** (1996), 225–242.

[34] H.J. Ryser, Combinatorial properties of matrices of zeros and ones, *Canadian Journal of Mathematics* **9** (1957), 371–377.

[35] A. Sinclair, Improved bounds for mixing rates of Markov chains and multicommodity flow, *Combinatorics, Probability and Computing* **1** (1992), 351–370.

[36] A. Steger and N. Wormald, Generating random regular graphs quickly, *Combinatorics, Probability and Computing* **8** (1999), 377–396.

[37] R. Taylor, Constrained switching in graphs, in *Combinatorial Mathematics VIII*, Springer Lecture Notes in Mathematics vol. 884, 1981, pp.314–336.

[38] W.T. Tutte, A short proof of the factor theorem for finite graphs, *Canadian Journal of Mathematics* **6** (1954), 347–352.

[39] R. Tyshkevich, Canonical decomposition of a graph (in Russian), *Doklady Akademii Nauk BSSR* **24(8)** (1980), 677–679.

[40] R. Tyshkevich, Decomposition of graphical sequences and unigraphs, *Discrete Mathematics* **220** (2000), 201–238.

[41] J.Y. Zhao, Expand and Contract: Sampling graphs with given degrees and other combinatorial families, Preprint, 2013. arXiv:1308.6627